MICROMACHINED STIMULATING ELECTRODES

Quarterly Report #2

(Contract NIH-NINDS-N01-NS-5-2335)

January 1996 --- April 1996

Submitted to the

Neural Prosthesis Program
National Institute of Neurological Disorders and Stroke
National Institutes of Health

by the

Center for Integrated Sensors and Circuits
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, Michigan
48109-2122

May 1996

MICROMACHINED STIMULATING ELECTRODES

Summary

During the past quarter, work on the development of stimulating electrode arrays for use in neural prostheses went forward in several areas. Equipment has been assembled that now allows the use of electrochemical impedance spectroscopy (EIS) and cyclic voltammetry (CV) to evaluate chronically-stimulated electrodes in-vitro or in-vivo over expanded ranges. EIS spectra can be taken from 10Hz to 1MHz, while CVs can be collected at scan rates from 50mV/sec to 1000mV/sec. A potentiostat (EG&G#283) generates CV waveforms and collects the data under GPIB control using PC-based LabVIEW software (National Instruments). EIS is performed using the same potentiostat in tandem with a frequency response analyzer (EG&G#1025). EIS experiments are also controlled via GPIB using special EIS software (EG&G#398). In-vivo data is reported from a chronic 1600μm^2 site implanted for 10 days both before and after the stimulation protocol. This testing capability should allow us to better understand electrode behavior in-vivo and changes that occur as the result of stimulation.

Several design modifications for the STIM-2 active probe have been considered during the past quarter. Alternative DAC designs have been explored and the problem with a weak pull-down of the current sinking function has been traced to insufficient bias current combined with the rather large distributed capacitance of the replicated transistors developing the output sink current for the DAC. This is being corrected by increasing the bias current for the DAC and, possibly, via a change in the circuitry for this function. The design of a reduced-complexity 64-site 4-channel active probe (STIM-2b) has now also nearly been completed. This probe steers the current generated on four external lines to four sites selected from 64 on the probe. The user can either stimulate or record from any of the selected sites as desired. A fabrication run of the monopolar (STIM-1b) and bipolar (STIM-1a) probes is also in process and should be completed during the coming term. Considerable study has been given to the problem of forming reliable contacts on these probes using dry (RIE) etching technology, and a number of processing options to ensure low contact resistance are being pursued. Finally, a number of design options for a second-generation external interface to the active probes are being explored. Five of the first-generation external interfaces have been completed and are ready for use with the probes.

MICROMACHINED STIMULATING ELECTRODES

1. Introduction

The goal of this research is the development of active multichannel arrays of stimulating electrodes suitable for studies of neural information processing at the cellular level and for a variety of closed-loop neural prostheses. The probes should be able to enter neural tissue with minimal disturbance to the neural networks there and deliver highly-controlled (spatially and temporally) charge waveforms to the tissue on a chronic basis. The probes consist of several thin-film conductors supported on a micromachined silicon substrate and insulated from it and from the surrounding electrolyte by silicon dioxide and silicon nitride dielectric films. The stimulating sites are activated iridium, defined photolithographically using a lift-off process. Passive probes having a variety of site sizes and shank configurations have been fabricated successfully and distributed to a number of research organizations nationally for evaluation in many different research preparations. For chronic use, the biggest problem associated with these passive probes concerns their leads, which must interface the probe to the outside world. Even using silicon-substrate ribbon cables, the number of allowable interconnects is necessarily limited, and yet a great many stimulating sites are ultimately desirable in order to achieve high spatial localization of the stimulus currents.

The integration of signal processing electronics on the rear of the probe substrate (creating an "active" probe) allows the use of serial digital input data which can be demultiplexed on the probe to provide access to a large number of stimulating sites. Our goal in this area has been to develop a family of active probes capable of chronic implantation in tissue. For such probes, the digital input data must be translated on the probe into per-channel current amplitudes which are then applied to the tissue through the sites. Such probes require five external leads, virtually independent of the number of sites used. As discussed in our previous reports, we are now developing a series of active probes containing CMOS signal processing electronics. Two of these probes are slightly redesigned versions of an earlier first-generation set of designs and are designated as STIM-1a and STIM-1b. The third probe, STIM-2, is a second-generation version of our high-end first-generation design, STIM-1. All three probes provide 8-bit resolution in setting the per-channel current amplitudes. STIM-1A and -1B offer a biphasic range using ±5V supplies from 0μA to ±254μA with a resolution of 2μA, while STIM-2 has a range from 0 to ±127μA with a resolution of 1μA. STIM-2 offers the ability to select 8 of 64 electrode sites and to drive these sites independently and in parallel, while -1a allows only 2 of 16 sites to be active at a time (bipolar operation). STIM-1b is a monopolar probe, which allows the user to guide an externally-provided current to any one of 16 sites as selected by the digital input address. The high-end STIM-2 contains provisions for numerous safety checks and for features such as remote impedance testing in addition to its normal operating modes. It also offers the option of being able to record from any one of the selected sites in addition to stimulation.

During the past quarter, we have continued to fabricate passive probe structures for internal and external users. No further problems with iridium adhesion have been experienced. New equipment has been installed for doing electrode impedance spectroscopy and cyclic voltammetry both in-vitro and in-vivo, and the first sites have been characterized before and after chronic stimulation. Progress has been made on the redesign of the high-end active probe, STIM-2, and the design of a new reduced-complexity probe,

STIM-2b, has nearly been completed. Five copies of the external interface for the active probes have also been completed. The results in each of these areas are described below.

2. *In-Vivo Stimulation with Passive Probes*

In Vivo Iridium Characterization Studies

Progress was made in several areas in animal experiments this past quarter. Electrochemical instrumentation was received and implemented . This new equipment will allow more accurate impedance and cyclic voltammetry measurements to be made. Also, acute cochlear nucleus stimulation experiments were performed as a precursor to chronic experiments.

Two electrode analysis methods, electrochemical impedance spectroscopy (EIS) and cyclic voltammetry (CV), have been used to evaluate a chronically stimulated electrode. Previous experiments measured impedance at 1 kHz and charge storage at a scan rate of 500 mV/sec. Equipment received this past quarter will allow for more thorough testing. EIS spectra will range from 10 Hz to 1 MHz. CVs will be collected at scan rates of 50 mV/sec to 1000 mV/sec. The test setup is shown schematically in Fig. 1. A potentiostat (EG&G 283) generates CV waveforms and collects data under GPIB control using PC-based LabVIEW software (National Instruments). EIS is performed using the same potentiostat in tandem with a frequency response analyzer (EG&G 1025). EIS experiments were also controlled via GPIB by special EIS software (EG&G 398). This system will be used for *in vitro* and *in vivo* testing.

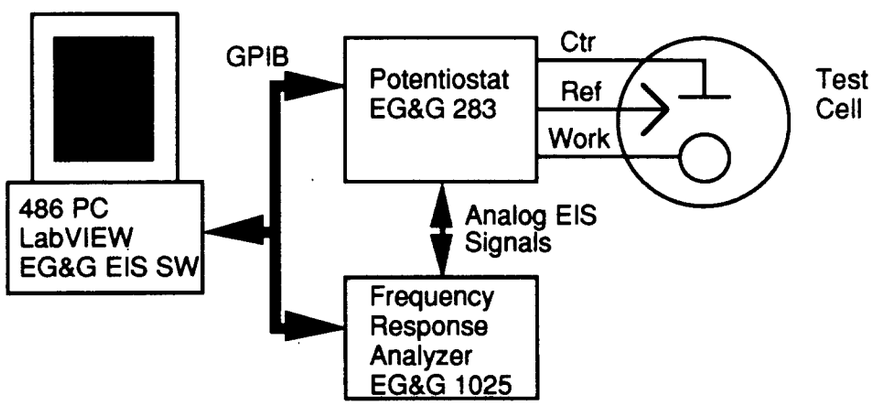

Fig. 1: Diagram of System for EIS and CV testing.

An initial experiment using the equipment above was completed this quarter. A chronic stimulating electrode was implanted in the cortex of a guinea pig. After a 10 day recovery period, monopolar stimulation through a 1600μm^2 site was performed. Biphasic current pulses (50μA, 100μsec/phase, 250 pulses/sec) were applied for four hours on five consecutive days. EIS and CV data were collected during the week of stimulation. The test cell consisted of the activated iridium stimulation site (working) and two 316 stainless steel head screws (counter and reference). As this was the first experiment using the new equipment, the testing did not proceed without problems, but some data was collected. EIS data from before and after the stimulation period on day 2 is shown in Fig. 2. The magnitude and phase spectra agree with the circuit model of the interface (parallel resistor

and capacitor). A drop in the impedance after stimulation is evident for mid-range frequencies. This result is consistent with 1 kHz measurements discussed in earlier reports and with other EIS data sets taken during this experiment. Those data showed a decline in impedance during the course of 4 hours of stimulation.

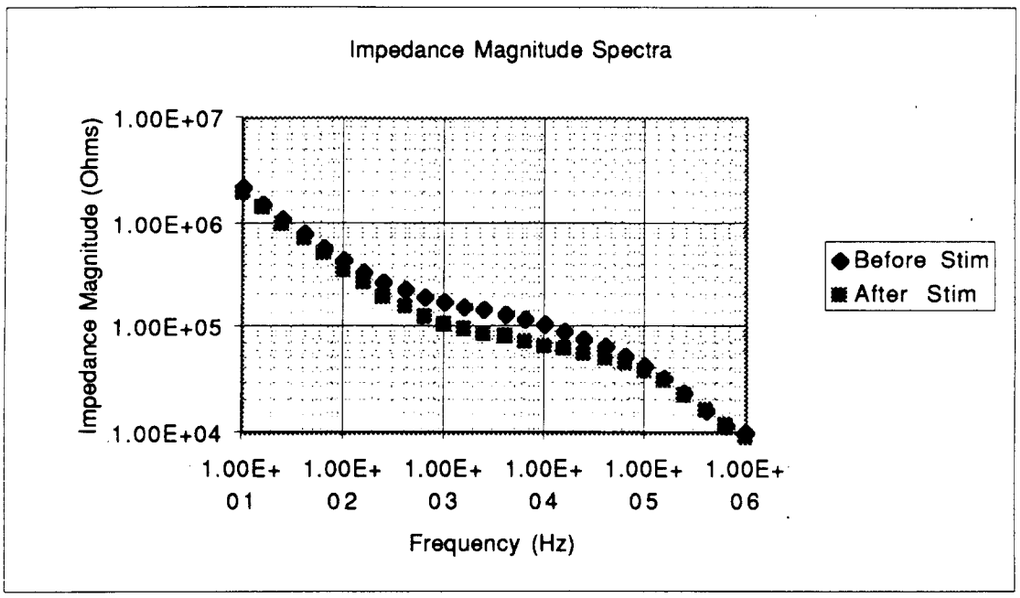

Fig. 2: Magnitude Spectra of a 1600 μm^2 site before and after 4 hours of stimulation.

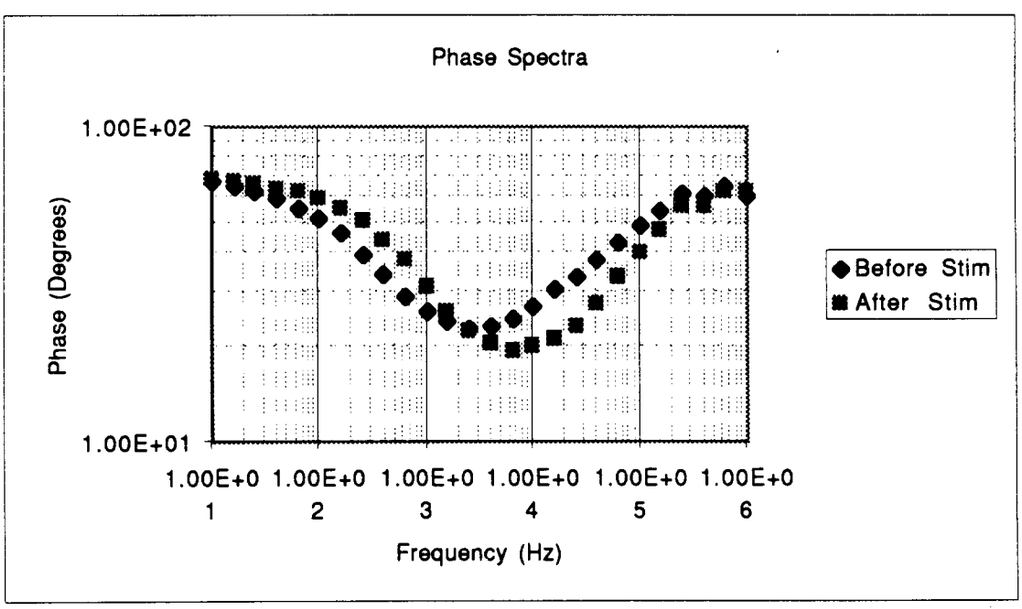

Fig. 3: Phase spectra of a 1600μm^2 site before and after 4 hours of stimulation.

Figure 4 shows CV data recorded on day 4 at scan rates of 100, 250, 500, and 1000 mV/sec. As expected, the current peaks grew with the scan rate. The integrated charge storage remained constant with increasing scan rate. No consistent changes were seen in CV data as a result of stimulation.

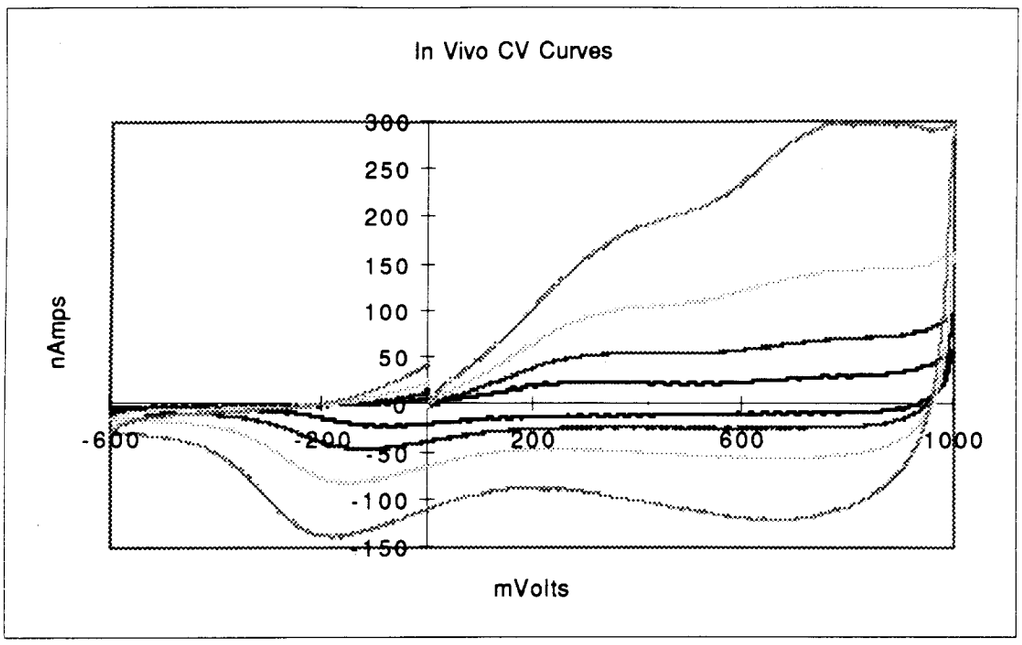

Fig. 4: In Vivo CV curves on 1600μm^2 sites at various scan rates.

The initial data show the new equipment will yield valuable information about the electrode-tissue interface. We plan to consult with EIC during the next quarter for assistance in analyzing our data and refining our testing protocol.

Several unsuccessful attempts were made to chronically implant a stimulating electrode in the cochlear nucleus of guinea pig. It is hoped that with a CN implant we can monitor evoked responses as well as impedance and possibly correlate changes in the threshold with changes in impedance spectra. We were able to record a middle latency response in an acute experiment (Fig. 5). The large peak of the electrical and acoustic response have similar shapes and, as expected, the electrical response has a shorter latency. The threshold for stimulation in this experiment was a 50μA, 100μsec biphasic pulse through a 2000μm^2 site. HMRI reports a much lower threshold (less than 10μA) with their CN stimulation. In the next quarter, we will perform several more acute experiments to practice our chronic surgical approach. If these are successful, we will again attempt a chronic implant.

During the next quarter we also hope to complete several chronic implant experiments. Histology should be available on the guinea pig discussed above. Acute CN stimulation experiments will be accomplished which will lead to chronic CN implants.

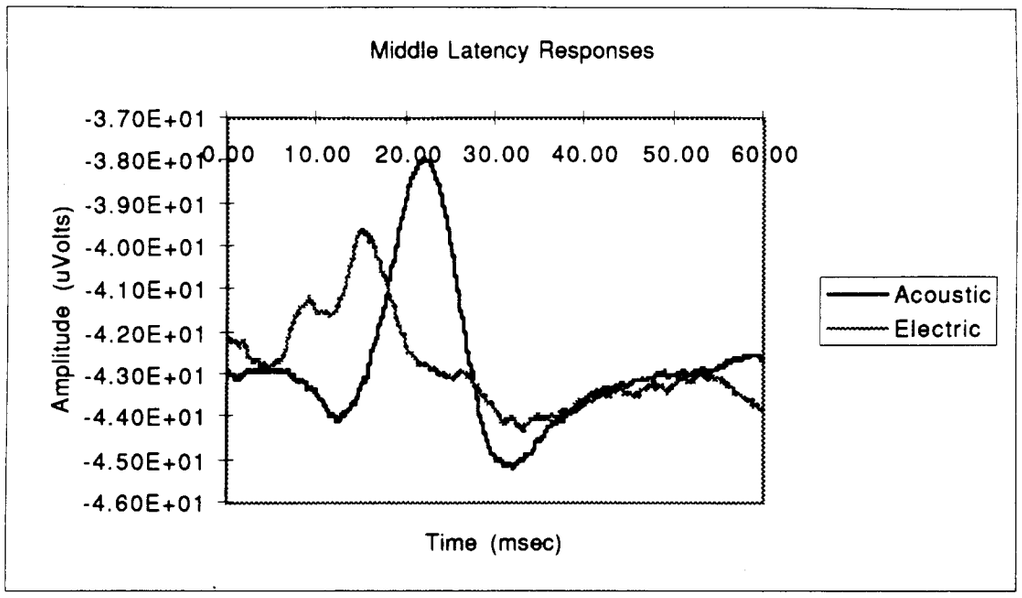

Fig. 5: Middle Latency Responses from Acute CN stimulation.

3. *Active Stimulating Probe Development*

During the past quarter, work on the active stimulating probes included the evaluation and simulation of various STIM-2 output circuit (DAC) designs, the design and simulation of the new STIM-2b probes, and the evaluation of various methods of improving our current method of making circuit contacts.

As noted in our earlier reports, the monopolar (STIM-1b) and bipolar (STIM-1a) versions of our active stimulating probes function well. New runs of these devices are underway, with additional in-vivo testing scheduled for the coming quarter. STIM-2, the high-end member of our present stimulating probe family, is currently being redesigned to correct several problems that caused yield problems in earlier process runs. This report will discuss some of these changes.

In the last quarterly report, several possible modifications of the STIM-2 DAC bias string were discussed. Further evaluation and simulation of the different bias string configurations have led to the following conclusions. For easy reference, the various circuits are included in this report; the current DAC biasing scheme is shown in Fig. 6. In this circuitry, bias currents are gated on when the channel is active and are mirrored to output sourcing (I_{pref}) or sinking (I_{nref}) drivers according to the desired current level. The first alternative design, Fig. 7, which utilizes secondary reference currents for independent control of the sourcing and sinking portions of the DAC, was discarded because of the large amount of circuit area it requires and its susceptibility to current mismatch due to being mirrored twice between the reference current and the output current. A second design, Fig. 8, which uses a bypass string for on/off control of the sourcing/sinking functions, was also discarded because of its susceptibility to device mismatch and associated changes in the reference current, i.e., any mismatch in the bias string and bypass string devices will be reflected by a change in the reference current, depending on which circuit branch the current is flowing through. The final design considered, Fig. 9, makes

use of a switch in the mirror path to realize control of the sourcing and sinking sections of the DAC. Both sections of the DAC mirror the same reference current which remains constant except during switching transients. For these reasons, this circuit is the configuration of choice.

The simulations of the STIM-2 DAC (Fig. 6) also revealed the answer to another problem that had been noticed with the DAC output: the slow turn-on of the current sink side of the DAC, as demonstrated very slow pull-down for sinking current in the oscilloscope trace of Fig. 10. A transient response simulation of the full DAC with all 127 unit current output transistors tied to each mirror bias point is shown in Fig. 11. An enlargement of just the DAC output current is shown in Fig. 12. The same slow turn-on (pull down) is observed. This is due to the very small (subnanoampere) sink bias current, which must charge up a very large capacitance associated with the 127 transistor gates until it reaches the output current transistor threshold voltage. The source bias current is about twice as large as the sink bias current as can be seen by the 3rd and 4th traces of Fig. 11. This would explain the difference in turn-on times, which are relatively slow, even for the sourcing devices. The current mirror bias voltages, as seen in the 1st and 2nd traces of Fig. 11, also demonstrate the problem by the difference in voltage ramp-up of the output transistor gates, which is then reflected in the current turn-on. The STIM-2 redesign includes a significant increase in the bias current, which will correct this problem by increasing the charge-up rate.

Another project during the past term has been the design of STIM-2b. STIM-2b is a second generation four-channel, 64-site version of the simplest active stimulating probe, STIM-1b. The probe uses a 20b shift register to load four 4b addresses, which are decoded to select one of sixteen sites for each channel. The selected site is connected to the analog data input/output pad through a pass-gate transistor to allow externally generated currents to be 'steered' to the selected site. A flag bit is included with the channel address in order to select between stimulation and a newly added recording function. The flag bit selects either a direct path between the I/O (current-input) pad and the site or selects the path through an output buffer amplifier for recording. A power-on-reset circuit (POR) sets all of the circuits to a startup state and connects all of the sites to the I/O pads in order to facilitate activation of all of the sites in parallel. The first clock pulse knocks down the POR signal and the probe subsequently begins functioning normally. The POR state can also be initiated by strobing the clock line, CSTB, low. A strobe on the address line, ASTB, resets all of the addresses to site zero and clears the flags (stimulation mode). Additional circuitry is included with site zero so that if the address line is held low, ASTB is maintained and all sites are disconnected, thereby allowing any probe leakage current to be measured. The functionality of this device is dependent mainly on the correct operation of the digital logic circuitry for selecting the correct site for current delivery or recording. The recording buffer amplifiers are essentially the only analog circuits, and their unity gain should make them very predictable and immune to variations in process parameters. The minimal amount of analog circuitry makes the design very robust with regard to its dependence on process parameters.

During the past quarter, the design of most of the STIM-2b circuit blocks has been completed. Circuit blocks from STIM-1b were used in this design whenever possible because of their proven functionality. A few circuit blocks from STIM-2 were also used directly or altered as necessary. Final optimization of transistor sizing, and so forth, will be done when the layout has been finished and the parasitic capacitances can be more accurately extracted (via layout software) and included in the simulations. Initial hand calculated estimates of these parasitics were extracted from the layouts of STIM-1b and STIM-2 and used in the initial simulations of circuit functionality and speed. These

preliminary simulations indicate that this probe should work at clock speeds as high as those attained by STIM-1b, above 4 MHz and possibly as high as 10 MHz.

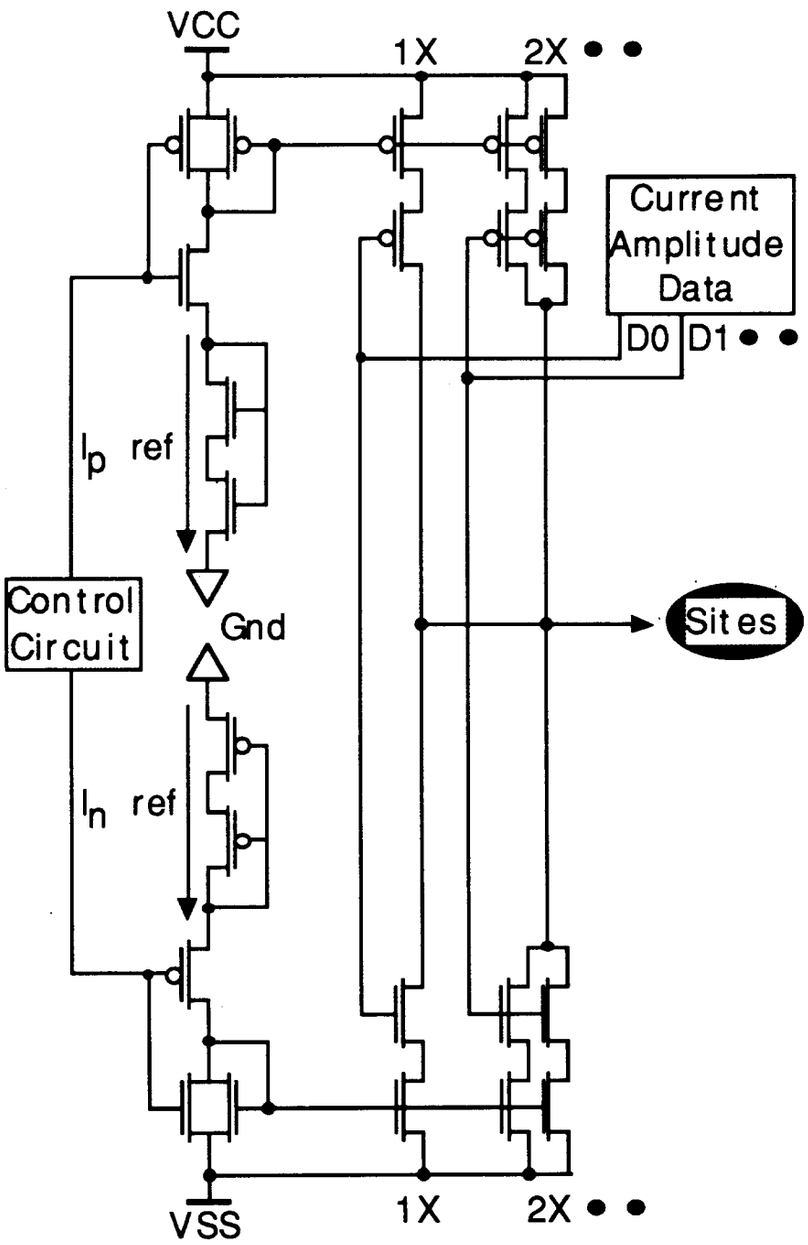

Fig. 6: Schematic of the present STIM-2 output DAC design.

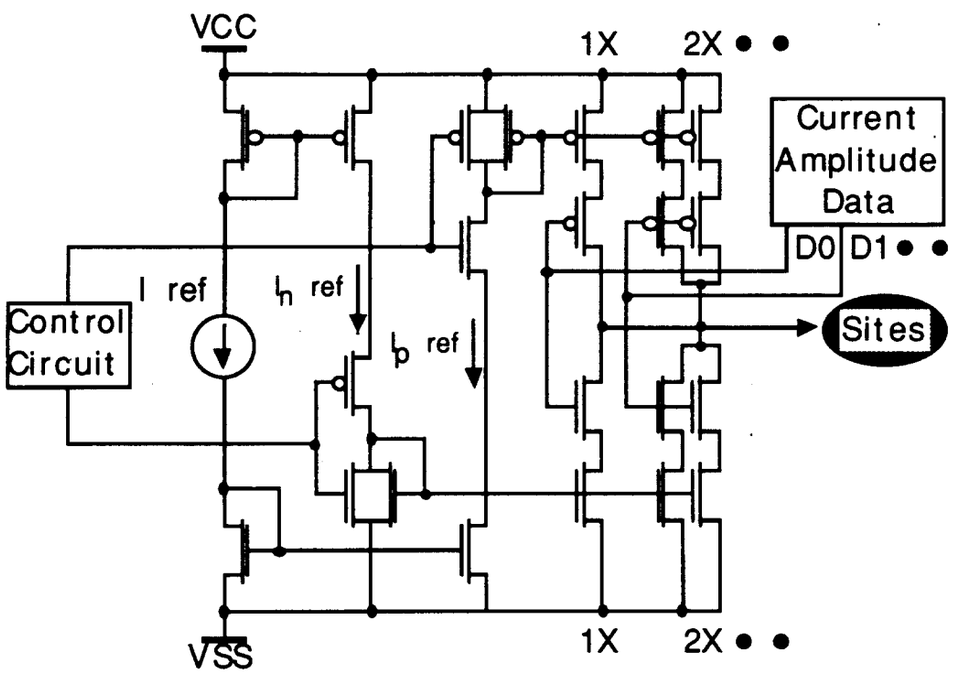

Fig. 7: Schematic of a circuit evaluated for a future implementation of the STIM-2 DAC. The design utilizes secondary reference currents for the sourcing and sinking portions of the DAC which can be independently switched on and off.

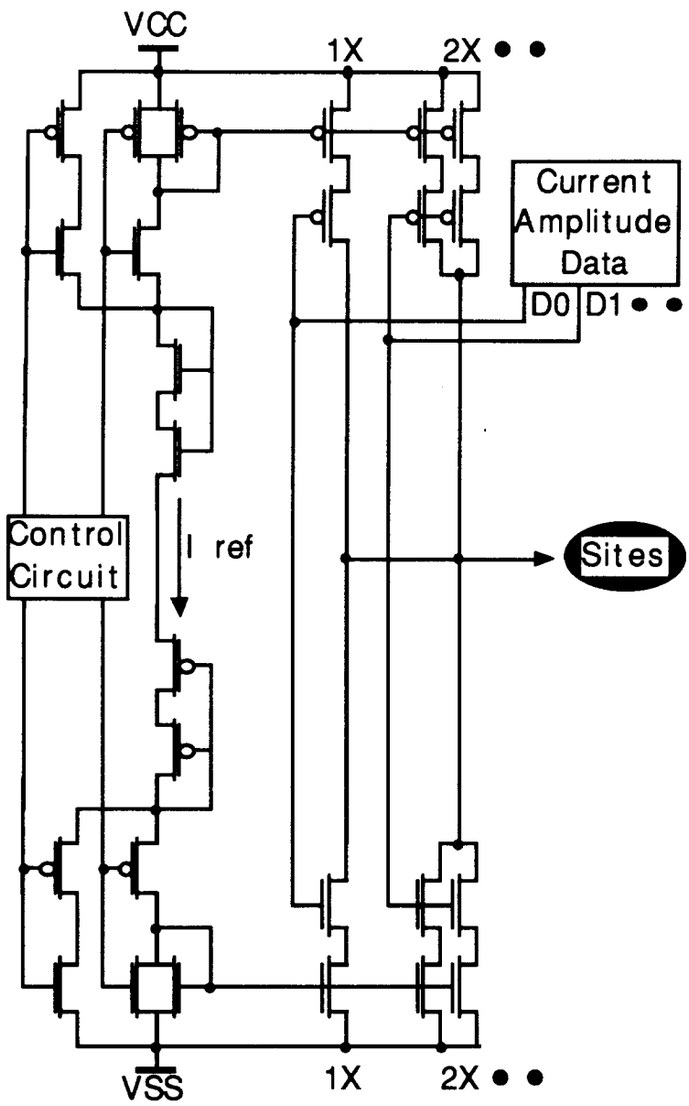

Fig. 8: Schematic of a second possible implementation of the STIM-2 DAC utilizing bypass strings for on/off control of the sourcing/sinking functions.

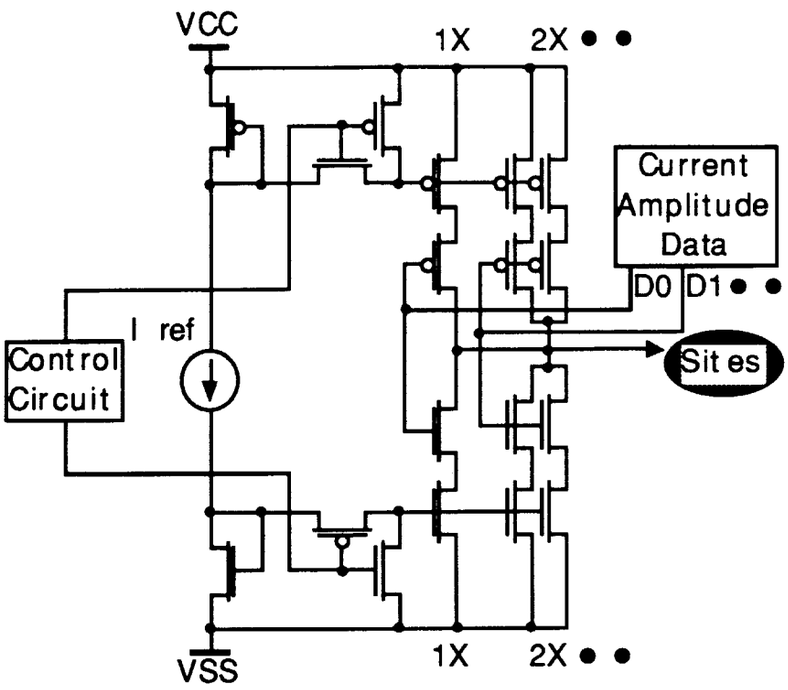

Fig. 9: Schematic of a third possible implementation of the STIM-2 DAC, which uses a switch in the mirror path to realize on/off control of the sourcing/sinking sections of the DAC.

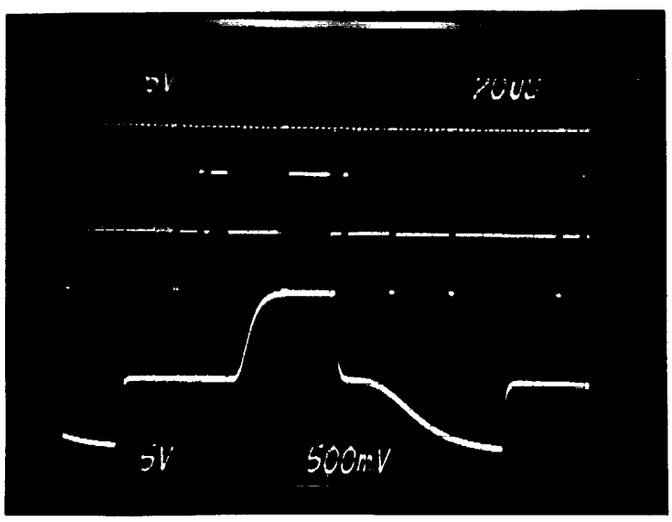

Fig. 10: The output of STIM-2, which demonstrates the slow turn-on of the DAC current sink even with the negative supply voltage turned up to increase the sink current.

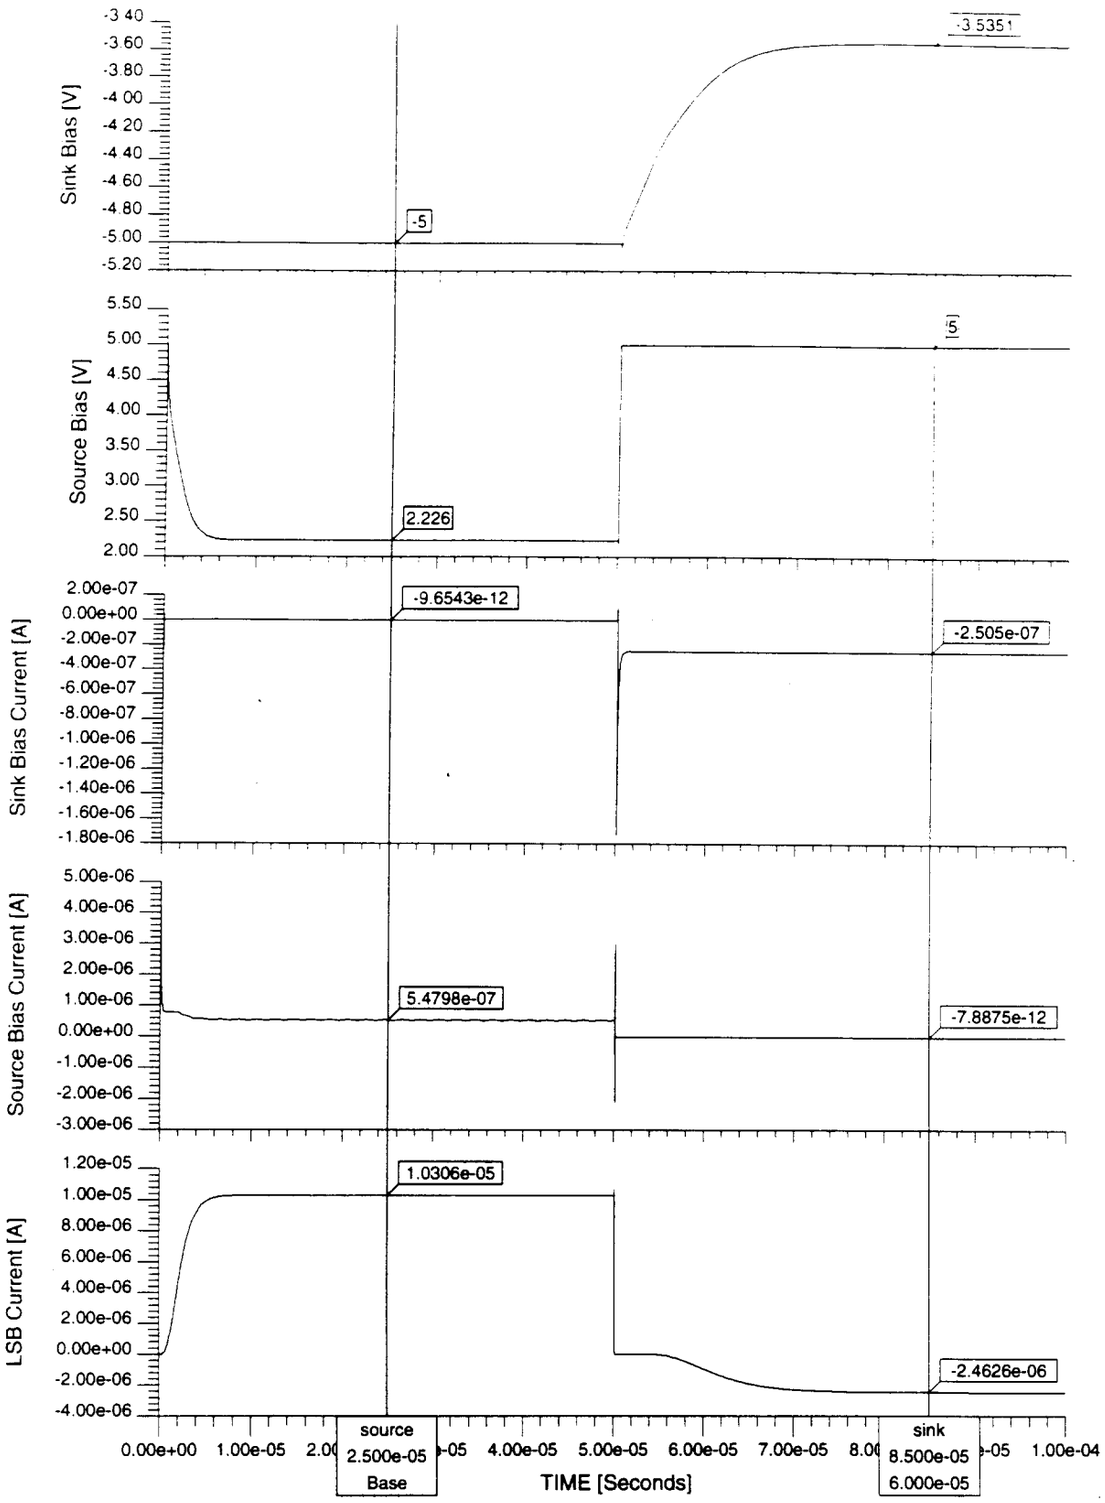

Fig. 11: The simulated output of STIM-2, which demonstrates the slow turn-on of the DAC current sink. The traces are (from top to bottom): the gate bias voltage on the current sink side of the DAC, the gate bias voltage on the current source side of the DAC, the bias string current on the current sink side of the DAC, the bias string current on the current source side of the DAC, and the DAC output current.

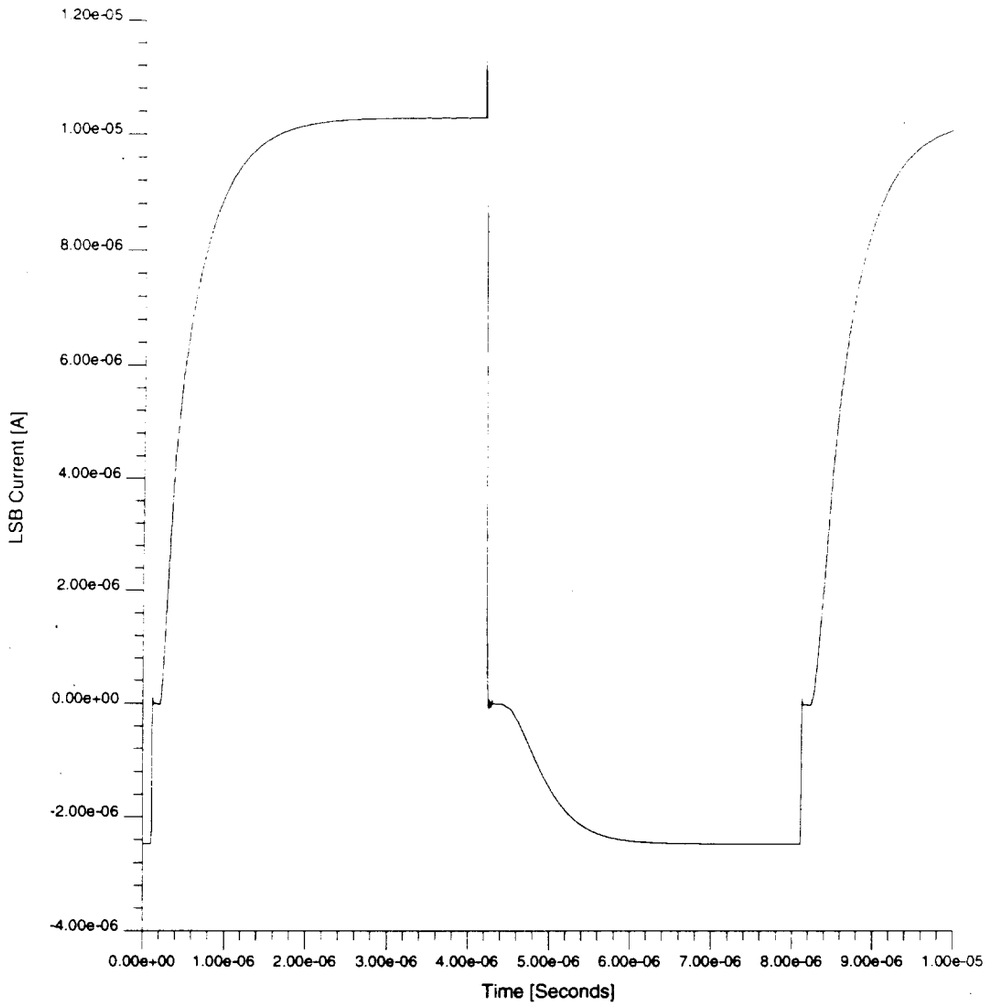

Fig. 12: An enlargement of the DAC output current, demonstrating the slow turn-on response of the current sink section.

Circuit Contacts

In the past, circuit contacts have caused significant problems in realizing working and reliable circuits. This has been attributed to various causes, such as polymer formation in the contact, low substrate surface doping levels (contacts not ohmic), and possible surface damage; many of these problems can be associated with the dry etching process being used, reactive ion etching (RIE). Attempts have been made to circumvent the problems by using an O_2 plasma etch in the RIE after the normal oxide etch, buffered HF dips prior to metal deposition, or annealing steps after contact formation. These have, in general, all met with some degree of success, but not with the reliability necessary to achieve yields as high as are desirable. During the past quarter, an extensive literature search was performed relative to the selective reactive ion etching of SiO_2-on-Si. The literature search for the most part simply confirmed that the postulated conditions can cause the contact problems being experienced; however, additional things that can cause bad contacts were also identified. The most important results were new suggested methods for forming good contacts. A short discussion of some of these results is summarized here.

Selective reactive ion etching of SiO_2-on-Si is generally based on fluorocarbon (freon) plasmas, as is the etch chemistry we use here in our lab. The formation of a polymer is inherent to this etch chemistry and is necessary for selectivity between SiO_2 and Si. The selectivity of the etch is based on the preferential formation of a polymer on Si and not on SiO_2. As the etch proceeds downward through the SiO_2 layer, very little polymer is formed on the SiO_2 allowing it to etch, but upon reaching the Si substrate or polysilicon, a substantial polymer layer begins to form, retarding the etching. This polymer remains on the substrate surface after a highly selective SiO_2-on-Si etch. Other problems such as deep charge traps due to lattice distortion (as deep as 30nm) can be caused by highly energetic photons emerging from the RIE plasma zone. Carbon, fluorine and hydrogen are implanted into the Si substrate due to bombardment during the etch, causing deep traps. The hydrogen can cause deactivation of boron doping by 'poisoning' or bonding to the acceptor site. A diagram of the layers present after a selective silicon dioxide etch are shown in Fig. 13.

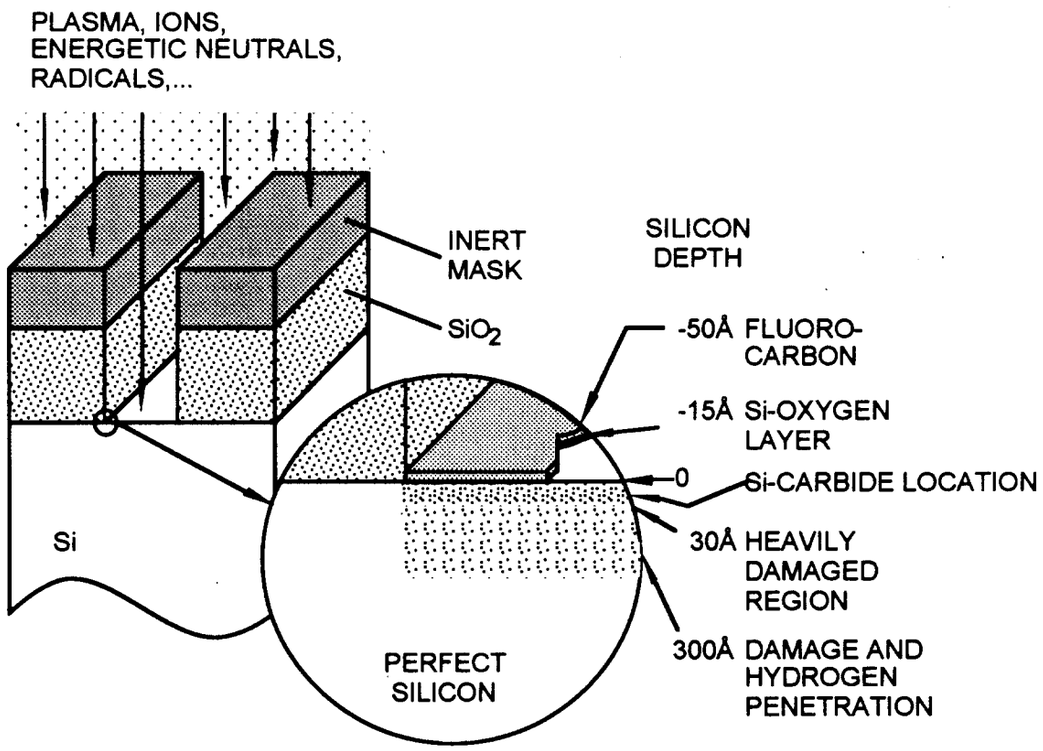

Fig. 13: Schematic view of changes of the near-surface properties of the silicon substrate resulting from a SiO2-on-Si selective reactive ion etch process.

In general, it appears that the polymer film and most of the carbon implanted in the amorphized top Si layer and most of the carbon implanted in the amorphized layer can be removed in an O_2 plasma. Lattice damage can normally be annealed out by a subsequent heating of the substrate at 400 to 800°C. It is necessary to remove the carbon before the high-temperature treatment in order to prevent the formation of silicon carbide at the silicon surface. Actually, annealing is seldom necessary if the silicon surface is to be used to make contact with a metal interconnect layer. Removal of the teflon-like layer is necessary and usually sufficient to ensure acceptable contact resistance, although it is not always possible

with an O_2 plasma etch. Any interfacial oxide produced in the O_2 plasma can be removed with HF if necessary. However, the problem with an O_2 plasma etch is that if performed in the same chamber in which the SiO_2 etch was performed, the polymer that is formed on the chamber walls (and it forms on the chamber walls during etching) is volitalized by the O_2 plasma, releasing free atomic fluorine which can then etch the Si substrate.

Another method for cleaning up the silicon surface is the treatment of the surface with a NF_3 plasma followed by a wet etch. This method has been shown to remove fluorocarbon films after as little as 10 seconds of exposure. Still another method is by using UV/O_3 to break up the polymer. Two exposures to UV/O_3 followed by HF dips are required to achieve nearly complete restoration of the Si surface. Unfortunately, these two procedures involve processes which are not immediately available to us.

As mentioned in the previous report, we have been evaluating two other methods of surface contamination removal: the use of chemical wet etchants, and the use of an oxygen anneal. The results are not yet conclusive insofar as the use of wet etching is concerned. A 400°C anneal in an O_2 ambient for 30 minutes can be used followed by a 30 second buffered HF dip to remove the oxide formed. This effectively breaks down the polymer film, anneals out lattice damage, and removes implanted contaminants with the oxide growth and removal. This does consume some silicon but it is small enough as to be insignificant.

During the coming quarter, we plan to perform a complete evaluation of the 400°C O_2 anneal and buffered HF dip as well as the use of chemical descum agents as possible post-etch treatment methods to achieve good circuit contacts. We also plan to finish the design, layout and simulation of the STIM-2b probe and will complete the current CMOS run of active stimulating probes following completion of the contact post-etch treatment evaluation.

4. External Probe Interface Development

During the past quarter we have received the final shipment of external electronic hardware assemblies. These consist of 5 units of the Chimera plug-in daughterboard and 5 units of the standalone remote converter unit. These assemblies were tested for errors in functionality that may have been introduced by the schematic capture, PCB layout, or manufacturing processes. Fortunately, no such errors were found and the manufactured assemblies behaved identically to the prototype units that served as the model for these designs. This completes the hardware development activity for the current generation of external interface systems.

In the software development effort, the graphical user interface for the external hardware remains stable and ready for use. This software consists of a Windows-hosted program that provides full control over the generation of pulsatile waveforms, plus a Chimera-hosted driver that accepts commands from the Windows program and translates them into the actual stimulation sequences. No further development of this system is expected in the short term, although it is anticipated that user-contributed suggestions will need to be implemented once the software enters into usage.

The focus of our efforts is shifting now towards the next generation of external interface systems. The main goal of this next stage of development is to achieve much greater flexibility in the bit-level communication protocol between the external interface and the probes themselves. The protocol of the current systems is hard-wired to support only

STIM-2, STIM-1, and STIM-1a probes. Any new probe design must either remain compatible with this protocol or we must adapt the external hardware for each new probe design. The preferred solution is to develop a hardware interface that is easily programmable and reconfigurable to support any protocol. This suggests that the bit-level waveform generation be placed under software control, rather than under a hardwired logic circuit.

Various architectures for the next generation of external interface systems are being considered. The architecture that is closest to the existing system is shown in Fig. 14 below. It simply does away with the daughterboard plug-in board for the Chimera and allows the DSP processor on the Chimera board to effect direct communication with the remote converter.

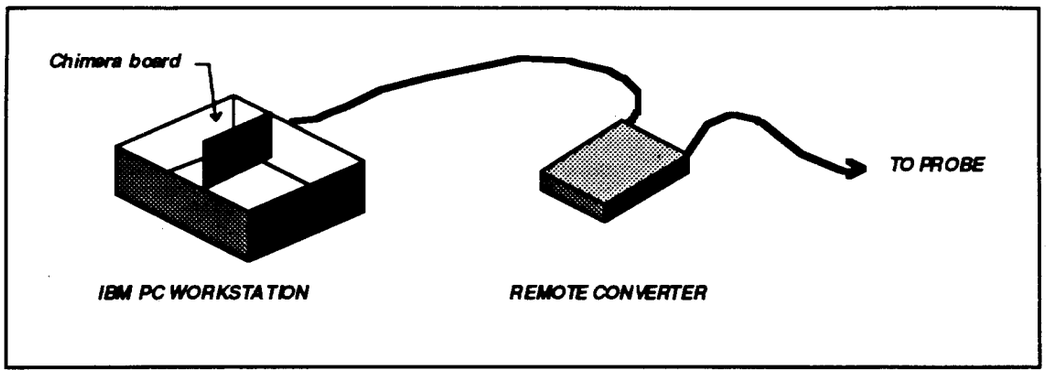

Fig. 14: A possible future implementation of the external probe electronic interface in which the DSP processor on the Chimera board directly drives the remote converter.

This architecture would eliminate the daughterboard from the system. Shifting the burden of bit-level waveform generation from the daughterboard to the Chimera board would allow for great flexibility in the protocol, due to the programmability of the Chimera's DSP processor, but may place a very large real-time burden on the Chimera, limiting the maximum bit rate that could be sustained.

One of the original reasons for choosing the Chimera board was its support for the plug-in daughterboard that is currently part of the system. Eliminating the daughterboard removes this advantage of the Chimera and allows us the freedom to use a different DSP board, perhaps one that is more powerful. Given that the Chimera was designed nearly a decade ago, it would be reasonable to assume that faster, more powerful DSP boards are on the market today and are available for comparable cost.

A different architecture eliminates the use of a PC plug-in board altogether and moves all of the waveform generation functions onto the remote converter. Thus, there would only be a single hardware unit responsible for all current external inteface functions. The remote converter would be linked to the PC via a standard serial connection. This architecture is diagrammed below in Fig. 15.

The advantages of this architecture include the fact that only a single hardware unit is required to implement the entire system and that the serial interface to the system allows any computer with a standard serial port to address the unit, not just an IBM PC. One disadvantage to this system is that the new hardware unit would require the design and

construction of new hardware, including some kind of CPU or DSP processor for the waveform generation and serial interface functions. This is likely to be an involved process. During the coming term, we will try to come to a conclusion concerning the preferred architecture for the next-generation system so that a more powerful interface can be realized.

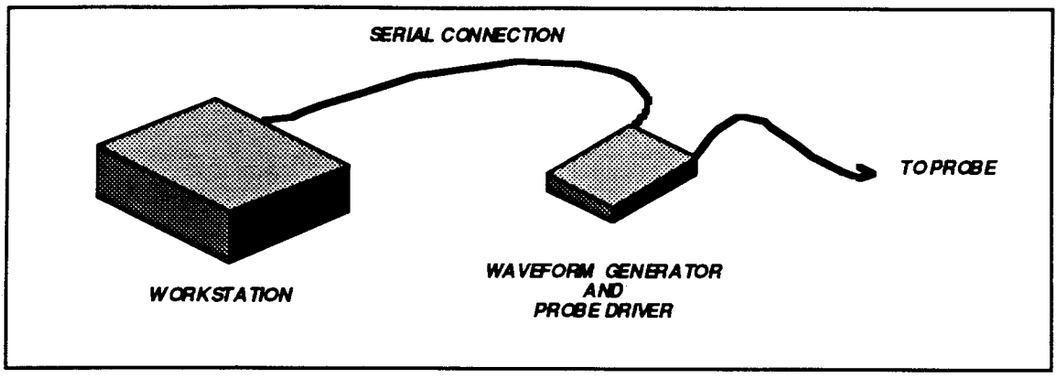

Fig. 15: External system architecture in which the desired waveform is completely generated in a remote converter.

5. *Conclusions*

During the past quarter, work on the development of stimulating electrode arrays for use in neural prostheses went forward in several areas. Equipment has been assembled that now allows the use of electrochemical impedance spectroscopy (EIS) and cyclic voltammetry (CV) to evaluate chronically-stimulated electrodes in-vitro or in-vivo over expanded ranges. EIS spectra can be taken from 10Hz to 1MHz, while CVs can be collected at scan rates from 50mV/sec to 1000mV/sec. A potentiostat (EG&G#283) generates CV waveforms and collects the data under GPIB control using PC-based LabVIEW software (National Instruments). EIS is performed using the same potentiostat in tandem with a frequency response analyzer (EG&G#1025). EIS experiments are also controlled via GPIB using special EIS software (EG&G#398). In-vivo data is reported from a chronic 1600μm^2 site implanted for 10 days both before and after the stimulation protocol. This testing capability should allow us to better understand electrode behavior in-vivo and changes that occur as the result of stimulation.

Several design modifications for the STIM-2 active probe have been considered during the past quarter. Alternative DAC designs have been explored and the problem with a weak pull-down of the current sinking function has been traced to insufficient bias current combined with the rather large distributed capacitance of the replicated transistors developing the output sink current of the DAC. This is being corrected by increasing the bias current for the DAC and, possibly, via a change in the circuitry for this function. The design of a reduced-complexity 64-site 4-channel active probe (STIM-2b) has now also nearly been completed. This probe steers the current generated on four external lines to four sites selected from 64 on the probe. The user can either stimulate or record from any of the selected sites as desired. A fabrication run of the monopolar (STIM-1b) and bipolar (STIM-1a) probes is also in process and should be completed during the coming term. Considerable study has been given to the problem of forming reliable contacts on these probe sing dry (RIE) etching technology, and a number of processing options to ensure

low contact resistance are being pursued. Finally, a number of design options for a second-generation external interface to the active probes are being explored. Five of the first-generation active interfaces have been completed and are ready for use with the probes.